

ABSTRACT

Method and apparatus for representing capacities and demands in a layered computing environment using normalized values for identifying optimal allocations of computing resources. A plurality of server models are established, with each server model including one or more server nodes and each server node having an associated set of capacity attributes. Similarly, a plurality of service models are established, each service model including one or more service nodes and each service node having an associated set of demand attributes. The server models are defined with a layered relationship as are the service models. A node that is part of a model in one layer corresponds to a model in the next-lower layer. The invention normalizes the demand and capacity attributes of the server nodes and service nodes in the different model layer in order to generate optimized mappings of service nodes to server nodes.